REMARKS

Claims 1-33 are pending in this application. Applicant respectfully requests reconsideration and allowance of the subject application.

35 U.S.C. § 103 – Saether and Strong

Claims 1-5, 10-12, 14-19, 22, 29 and 31-33 stand rejected under 35 U.S.C. §103(a) as being unpatentable over U.S. Patent No. 6,405,219 to Saether et al. (hereinafter "Saether") in view of U.S. Patent No. 5,689,688 to Strong et al. (hereinafter "Strong").

Claim 1 recites:

A method of synchronizing data among a plurality of web servers, wherein each of the plurality of web servers is coupled to a common data server, the method comprising:
retrieving a scheduled activation time from the data server;
prior to the scheduled activation time, retrieving updated data into staging caches in the plurality of web servers; and
at the scheduled activation time, copying the updated data from the staging caches within each of the plurality of web servers to an active cache within each of the plurality of web servers.

Thus, Claim 1 is directed at a method of synchronizing <u>data</u> among web servers. As recited in Claim 1, both a scheduled activation time and updated data are retrieved. By copying updated data from the staging caches to active caches within the web servers at the scheduled activation time, the web servers are concurrently synchronized. Neither Saether nor Strong discloses this method of synchronizing data.

Saether describes a technique for a global server to provide updated source files to content servers in a network. However, although Saether discloses updating files on a server, nothing in Saether discloses or suggests a method for synchronizing the files among multiple content servers using a scheduled activation time. The Office Action also acknowledges that Saether does not explicitly disclose retrieving a scheduled activation time from the data server. (See Office Action, page 2). Strong fails to cure these deficiencies.

Strong describes a method for synchronizing local times, maintained at nodes within a network architecture, with a reference time. (See Strong, Abstract). Although the method described by Strong may be used to synchronize local times at network nodes, this method does not synchronize the data on those nodes at a particular activation time. The Office Action appears to argue that the local times described in Strong are equivalent to the data in Claim 1. (See Office Action, page 11). However, the method recited in Claim 1 synchronizes data for web servers by copying updated data from a staging cache to an active cache in each web server at the scheduled activation time. Thus, the activation time and the update data are separate elements in Claim 1, and the local times in Strong are not equivalent to the updated data recited in Claim 1.

Nothing Saether or Strong suggests how the two references can be combined. The Office Action cites materials in Saether and Strong that allegedly provide the suggestion to combine. (See Office Action, page 11 and 12). The cited materials in Saether state:

Therefore, a need exists for a computer implementable method of distributing a set of current version of source files to a plurality of content servers using a minimal amount of bandwidth. (Saether, col. 1, lines 38-40)

This sentence only provides a motivation to implement the technique in Saether, and not to combine Saether's technique with the local time synchronization scheme described in Strong. The cited materials in Strong state:

> In order for the master node to synchronize other nodes, the master node conventionally must know which other nodes it is responsible for updating, so that it can direct appropriate time update messages to those nodes. (Strong, col. 2, lines 7-10)

This sentence only describes why the master node must know other nodes. However, it does not provide a motivation to combine the time synchronization scheme in Strong with the file distribution technique in Saether.

Even if Saether and Strong can be combined, the method described in Strong would merely allow the global server and the content server in Saether to operate in accordance with the same time reference. However, the combination of Saether and Strong still does not teach or suggest synchronizing <u>data</u> among web servers with <u>a scheduled activation time</u> retrieved from a data server, as recited in Claim 1.

For at least these reasons, Applicant respectfully submits that Claim 1 is allowable over Saether and Strong, alone or in combination. Given that Claims 2-5, 10-12 and 14 depend from Claim 1, Claims 2-5, 10-12 and 14 are also allowable for at least the same reasons.

Claim 15 recites:

> A system comprising:
> a plurality of web servers coupled to a common data server, wherein each of the plurality of web servers comprises:
>> a staging cache;
>> an active data cache coupled to the staging cache;
>> wherein the web server is configured to retrieve a scheduled activation time from the data server, and further configured to retrieve updated data from the data server into the staging cache prior to the scheduled activation time; and
>> wherein the web server is configured to copy data from the staging cache to the active data cache at the scheduled activation time.

As discussed above, neither Saether nor Strong discloses or suggests synchronizing data among multiple servers using a scheduled activation time. Saether and Strong also fail to disclose or suggest the web server recited in Claim 15, which is configured "to retrieve a scheduled activation time from the data server" and "to copy data from the staging cache to the active data cache at the scheduled activation time". Thus, for at least these reasons, Applicant respectfully submits that Claim 15 is allowable over Saether and Strong, alone or in combination. Given that Claims 16-19 and 22 depend from Claim 15, Claims 16-19 and 22 are also allowable for at least the same reasons.

Claim 29 recites:

> A method of synchronizing data among a plurality of web servers, wherein each of the plurality of web servers is coupled to a common data server, the method comprising:
>> providing a scheduled activation time from the data server to each of the plurality of web servers;
>> communicating updated data into a staging cache in each of the plurality of web servers prior to the scheduled activation time; and

copying data from the staging cache in each of the plurality of the web servers to an active cache in each of the plurality of the web servers at the scheduled activation time.

As discussed above, neither Saether nor Strong discloses or suggests synchronizing data among multiple servers using a scheduled activation time. Thus, Saether and Strong also fail to disclose or suggest "providing a scheduled activation time from the data server to each of the plurality of web servers" and "copying data from the staging cache in each of the plurality of the web servers to an active cache in each of the plurality of the web servers at the scheduled activation time", as recited in Claim 29.

Thus, for at least these reasons, Applicant respectfully submits that Claim 29 is allowable over Saether and Strong, alone or in combination. Given that Claims 31-33 depend from Claim 29, Claims 31-33 are also allowable for at least the same reasons.

35 U.S.C. § 103 – Saether, Strong and Hagersten

Claims 6 and 30 stand rejected under 35 U.S.C. §103(a) as being unpatentable over Saether in view of Strong and in further view of U.S. Patent No. 5,958,019 to Hagersten et. al (hereinafter "Hagersten").

Claims 6 and 30 recite:

> 6.   A method as recited in claim 1 wherein retrieving updated data into staging caches in the plurality of web servers is performed asynchronously.

30.   A method as recited in claim 29 wherein the communicating updated data into a staging cache is performed asynchronously.

As discussed above, neither Saether nor Strong discloses or suggests synchronizing data among multiple servers using a scheduled activation time, as recited in independent Claims 1 and 29. The Office Action also acknowledges that Saether, in view of Strong, does not explicitly disclose the subject matter in Claims 6 and 30. (See Office Action, page 5). Hagersten fails to remedy these deficiencies.

Hagersten teaches a multiprocessing system configured to perform synchronization operations. The Office Action asserts that Hagersten shows several examples of staging caches and active caches. (See Office Action, page 12). Although Hagersten describes various types of caches in a computer system, Hagersten fails to disclose or suggest asynchronously retrieving or communicating updated data into staging caches in servers, as recited in Claims 6 and 30. Therefore, the combination of Saether, Strong and Hagersten fails to disclose or suggest all elements of Claims 6 and 30.

For at least these additional reasons, Applicant respectfully submits that Claims 6 and 30 are allowable over Saether, Strong and Hagersten, alone or in combination.

35 U.S.C. § 103 – Saether, Strong and Yamazaki

Claims 7 and 20 stand rejected under 35 U.S.C. §103(a) as being unpatentable over Saether in view of Strong and in further view of U.S. Patent No. 5,923,855 to Yamazaki (hereinafter "Yamazaki").

Claims 7 and 20 recite:

7.   A method as recited in claim 1 further comprising:
after the scheduled activation time, updating data caches in the data server.

20.   A system as recited in claim 15 wherein each of the plurality of web servers is configured to update data caches in the data server after the scheduled activation time.

As discussed above, Saether and Strong fail to teach or suggest synchronizing data among multiple servers using a scheduled activation time. The Office Action also acknowledges that Saether, in view of Strong, does not explicitly disclose after the scheduled activation time, updating data caches in the data server. (See Office Action, page 6). Yamazaki does not cure these deficiencies.

Yamazaki discloses a multi-processing system and method for synchronizing among processors with cache memory having reset state, invalid state and valid state. Although Yamazaki describes updating data stored in cache memories of different processors in a multi-processor system, Yamazaki fails to disclose or suggest updating data caches after a scheduled activation time, as recited in Claims 7 and 20. Therefore, the combination of Saether, Strong and Yamazaki fails to disclose or suggest the elements of Claims 7 and 20. For at least these additional reasons, Applicant respectfully submits that Claims 7 and 20 are

allowable over Saether, Strong, and Yamazaki, alone or in combination.

35 U.S.C. § 103 – Saether, Strong and Sakon

Claims 8-9 and 21 stand rejected under 35 U.S.C. §103(a) as being unpatentable over Saether in view of Strong and in further view of U.S. Patent No. 5,796,946 to Sakon (hereinafter "Sakon").

Claims 8-9 and 21 recite:

> 8.   A method as recited in claim 1 further comprising:
> after the scheduled activation time, calculating a next scheduled activation time.

> 9.   A method as recited in claim 1 further comprising:
> after the scheduled activation time, updating data caches in the data server and calculating a next scheduled activation time, wherein the updating and calculating are performed by the first web server to initiate a retrieval process after the scheduled activation time.

> 21.   A system as recited in claim 15 wherein each of the plurality of web servers is configured to calculate a next scheduled activation time after the scheduled activation time.

As discussed above, neither Saether nor Strong discloses or suggests synchronizing data among multiple servers using a scheduled activation time. Thus, Saether and Strong fail to teach or suggest calculating such a time in a manner as recited in claims 8-9 and 21. The Office Action also acknowledges that Saether, in view of Strong, does not explicitly disclose calculating the next scheduled activation time. (See Office Action, page 7). Sakon does not remedy these deficiencies.

Sakon discloses multi-processor system barrier synchronizer. Although

Sakon describes a next barrier value, such a value is used to synchronize processors in a multi-processor system, not for synchronizing data among multiple servers. Thus, the next barrier value described by Sakon is not equivalent to the next scheduled activation time as recited in Claims 8-9 and 21. Therefore, the combination of Saether, Strong and Sakon fails to teach or suggest the limitations of Claims 8-9 and 21.

For at least these additional reasons, Applicant respectfully submits that Claims 8-9 and 21 are allowable over Saether, Strong and Sakon, alone or in combination.

35 U.S.C. § 103 – Saether, Strong and Brendel

Claim 13 stands rejected under 35 U.S.C. §103(a) as being unpatentable over Saether in view of Strong and in further view of U.S. Patent No. 5,774,660 to Brendel et al. (hereinafter "Brendel").

Claim 13 recites:

> A method as recited in claim 1 wherein the plurality of web servers comprise a web farm, and wherein the plurality of web servers are load balanced using a domain name service (DNS) round-robin technique.

As discussed above, neither Saether nor Strong discloses or suggests a method for synchronizing data among multiple servers using a scheduled activation time, which is the subject matter recited in Claim 1.

Brendel discloses a world-wide web server with delayed resource-binding for resource-based load balancing on a distributed resource multi-node network.

Although Brendel teaches the use of DNS, nothing in Brendel remedies the deficiencies of Saether and Strong discussed above. Thus, the combination of Saether, Strong and Brendel fails to teach or suggest the elements recited in Claim 13. For at least these reasons, Applicant respectfully submits that Claim 13 is allowable over Saether, Strong and Brendel.

35 U.S.C. § 103 – Saether, Strong, Yamazaki and Sakon

Claims 23-28 stand rejected under 35 U.S.C. §103(a) as being unpatentable over Saether in view of Strong in further view of Yamazaki in further view of Sakon.

Claim 23 recites:

> One or more computer-readable media having stored thereon a computer program comprising the following steps:
> retrieving a scheduled activation time from a data server;
> prior to the scheduled activation time, retrieving updated data into a staging cache in a server;
> at the scheduled activation time, copying data from the staging cache in the server to an active cache in the server; and
> after the scheduled activation time, updating data caches in the data server and calculating a next scheduled activation time.

As discussed above, nothing in Saether, Strong, Yamazaki, or Sakon discloses or suggests synchronizing data among multiple servers using a scheduled activation time, updating of data caches after a scheduled activation time, or calculating a next scheduled activation time. Thus, the combination of these references still fails to teach or suggest the subject matter of Claim 23. For at least these reasons, Applicant respectfully submits that Claim 23 is allowable over
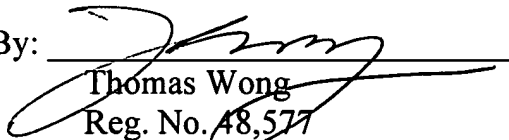
Saether, Strong, Yamazaki and Sakon, alone or in combination. Given that Claims 24-28 depend from Claim 23, Claims 24-28 are also allowable over Saether, Strong, Yamazaki and Sakon for at least the same reasons

Conclusion

Claims 1-33 are now in condition for allowance. Applicant respectfully requests the issuance of the subject application. Should any matter in this case remain unresolved, the undersigned attorney respectfully requests a telephone conference with the Examiner to resolve any such outstanding matter.

Respectfully Submitted,

Date: 8/10/2004

By: _____
Thomas Wong
Reg. No. 48,577
(206) 315-4001 x106